arXiv:1412.6170v1 [cs.DC] 18 Dec 2014

Manycore processing of repeated k-NN queries over massive moving objects observations

Francesco Lettich [*], Salvatore Orlando [†], and Claudio Silvestri [‡]

Dipartimento di Scienze Ambientali, Informatica e Statistica, Università Ca' Foscari
Via Torino 155, Venice, Italy

December 22, 2014

Abstract

The ability to timely process significant amounts of continuously updated spatial data is mandatory for an increasing number of applications. In this paper we focus on a specific data-intensive problem concerning the repeated processing of huge amounts of k nearest neighbours (k-NN) queries over massive sets of moving objects, where the spatial extents of queries and the position of objects are continuously modified over time. In particular, we propose a novel hybrid CPU/GPU pipeline that significantly accelerate query processing thanks to a combination of ad-hoc data structures and non-trivial memory access patterns.

To the best of our knowledge this is the first work that exploits GPUs to efficiently solve repeated k-NN queries over massive sets of continuously moving objects, even characterized by highly skewed spatial distributions. In comparison with state-of-the-art sequential CPU-based implementations, our method highlights significant speedups in the order of 10x-20x, depending on the datasets, even when considering cheap GPUs.

1 Introduction

An increasing amount of applications need to process massive spatial workloads. Specifically, we consider applications in settings where spatial data is continuously produced over time and needs to be processed rapidly, e.g., scenarios involving services or applications mainly tailored for mobile device infrastructures - such as *Location-Based Services* (LBS) or *Location-Based Social Networking* applications (LBSN) [1], and others, like behavioural simulations or Massively Multiplayer Online Games (MMOG) [2]. In these applications, very large populations of continuously *moving objects* frequently update their positions and issue some kind of query to look for other objects within their interaction area. The resulting massive workloads pose new challenges to data management techniques. In this context we consider k-NN queries, i.e., globular queries whose spatial extent is dictated by the dislocation of the k nearest objects with respect to their centers.

To enable parallel processing and optimizations, and thus manage the targeted workloads in a scalable manner, we recur to an approach based on time discretization. In this sense, we partition the time in intervals (or *ticks*), assign location updates and queries to the ticks in which they occur and process the updates and queries in the resulting batches such that query results are reported after the end of each tick. This approach has the effect of replacing the processing of a large number of independent and asynchronous queries with the *iterated processing of spatial joins* between the last known positions of all the moving objects at the end of a tick and the queries issued during the tick. In other words we trade (slightly) delayed processing of queries for increased throughput, and therefore care is needed to ensure acceptable delays. To achieve high performance and scalability we also exploit a platform encompassing an off-the-shelf general-purpose microprocessor (CPU) coupled with a *Graphics Processing Unit* (GPU) that features hundreds of processing cores. To benefit from GPUs exploitation, limitations and peculiarities of these architectures must be carefully taken into account [3]. Specifically, individual GPU cores are slower than those of a typical CPU, while some memory access patterns may cause serious performance

[*] lettich@dais.unive.it
[†] orlando@unive.it
[‡] silvestri@unive.it

degradation due to contention and serialization of memory accesses. Effective query processing techniques are needed to address these limitations: multiple cores must work together to efficiently process queries in parallel for most of the time, and must coordinate their activities to ensure high memory bandwidth.

Our contributions can be summarized as follows: first, we introduce a framework for repeated processing of massive k-NN queries over massive moving object observations. Second, on the basis of this framework we introduce K-NN$_{\text{GPU}}$, an hybrid CPU-GPU k-NN query processing pipeline able to efficiently compute batches of k-NN queries. K-NN$_{\text{GPU}}$ processes queries using an iterative approach, where at each iteration blocks of close objects are scanned in parallel to update the result sets of blocks of k-NN queries. At each iteration we exploit proper memory access patterns coupled with ad-hoc memory layouts of objects and queries, which are arranged according to some spatially preserving function: this entails workloads suitable for GPU processing and improves memory bandwidth due to better exploitation of GPUs caching and coalescing capabilities. Finally, we carry out an extensive set of experiments to study how K-NN$_{\text{GPU}}$ varies its performance for different datasets and algorithm parameters. We also compare K-NN$_{\text{GPU}}$ with a GPU baseline [4] and with a state-of-the-art, sequential CPU-based algorithm.

The paper is organized as follows: in Section 2 we provide the formal framework used to model the processing of k-NN queries. In Section 3 we give a brief overview about modern GPUs and review the main ideas behind K-NN$_{\text{GPU}}$, as well as the main data structures used, while Section 4 presents the k-NN query processing pipeline. Finally, Sections 5 and 6 provide the experimental part of our work, Section 7 covers the related work while Section 8 gives the conclusions. Finally, Appendix A provides further details about the complexity of the various pipeline phases.

2 Problem Setting and Statement

2.1 Problem Setting

We consider a set of points $O = \{o_1, \ldots, o_n\}$ moving in a two-dimensional Euclidean space $\mathbb{R}^2$, where the position of object o_i is given by the function $pos_i : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^2$ mapping time instants into spatial positions. These points model objects that issue position updates and queries as they move. Let $\mathcal{P}_i = \langle p_i^{t_0}, \ldots, p_i^{t_k}, \ldots \rangle$, $t_j < t_{j+1}$, be the time-ordered sequence of position updates issued by o_i, where $p_i^{t_j} = pos_i(t_j)$ is a position update. For a generic time t, $t \geq t_0$, the most recently known position of o_i *before* t is denoted by $\hat{p}_i^t$, and defined as follows:

$$\hat{p}_i^t = p_i^{t_m} \in \mathcal{P}_i \text{ if } t_m < t \leq t_{m+1}$$

The goal of a generic k-NN query $q_i = (x, y, k)$, is to find the k closest objects to its center (x, y). Since we are dealing with moving objects, and the processing time t' of the query may be later than its issuing time, we have to consider the most recently known positions of objects.

Definition 1 (K-NN QUERY) *Let $q_i = (x, y, k)$ be a k-NN query, issued by object o_i, centered in (x, y). Its result set, computed at time t, is defined as:*

$$res(q_i, t) = \{o_j \in O \mid rank(q_i, o_j, t) < k \ \wedge \ i \neq j\},$$

where

$$rank(q_i, o_j, t) = |\{o \in O \mid dist_t(q_i, o) < dist_t(q_i, o_j)\}| \ \geq \ 0$$

is the number of objects whose most recently known positions at time t are closer to the center of q_i than the position of the most recent update of o_j. Thus $dist_t(q_i, o_j)$ is the Euclidean distance between the query position and the last known position $\hat{p}_j^t$ of object o_j at time t.

According to the above definition we observe that, when there are ties in rank comparison, the result set of a query may contain more than k elements. When this happens, without loss of generality in this work we always return k elements for each query by arbitrarily selecting a subset of the elements having the same maximum rank.

2.2 Batch Processing

In this work we assume that the processing of queries can be delayed up to a certain extent to optimize the overall system throughput when facing intensive workloads deriving from massive numbers of moving objects and

queries, yet satisfying possible, specific QoS requirements. We quantize the time into *ticks* (time intervals) with the objective of processing updates and queries in batches at the end of each tick. Given a moving object, if we observe multiple position updates and queries issued during a time tick, only the most recent ones are processed. This procedure for computing queries ensures serializable query processing and implements the timeslice query semantics, where query results are consistent with the database state at a given time, usually when we start processing the query [5].

Hereinafter we will use the following notation: Δt is the tick duration, $\tau_k = [k \cdot \Delta t,\ \ (k+1) \cdot \Delta t)$ is the k-th tick, $P^{\tau_k} = \{p_1^{\tau_k}, \ldots, p_n^{\tau_k}\}$ is the set of last known positions of all objects at the beginning of the next $(k+1)$-th tick. Similarly, $Q^{\tau_k} = \{q_1^{\tau_k}, \ldots, q_n^{\tau_k}\}$ are the most recent queries issued during τ_k and known at the beginning of the next $(k+1)$-th tick ($q_i^{\tau_k} = \perp$ in case no query was issued by o_i during τ_k).

2.3 Problem Statement

Given *(i)* a set of n objects O, *(ii)* a partitioning of the time domain into ticks $[\tau_k]_{k \in \mathbb{N}}$ of duration Δt, *(iii)* a query latency requirement λ, and *(iv)* a sequence of pairs $[(P_{\tau_k}, Q_{\tau_k})]_{k \in \mathbb{N}}$, where P_{τ_k} is the set of up-to-date object positions at the end of τ_k, and Q_{τ_k} is the set of the last issued queries during τ_k, the goal of **iterated batch processing** of queries Q_{τ_k} over the corresponding P_{τ_k}, $k \in \mathbb{N}$, is to **compute** $[R_{\tau_k}]_{k \in \mathbb{N}}$, i.e., a **sequence of pairs**, each composed of a query and the list of the corresponding results:

$$R_{\tau_k} = \{(q_i^{\tau_k},\ res(q_i^{\tau_k},\ (k+1) \cdot \Delta t)) \mid q_i^{\tau_k} \neq \perp \wedge q_i^{\tau_k} \in Q_{\tau_k}\}$$

3 Proposal overview

The use of GPUs for general purpose computation presents significant advantages, despite the more complex computational model and memory architecture. In this paper we refer to NVIDIA CUDA framework and terminology, however noting that different frameworks and architectures adopt similar solutions under different names. Here we summarize some terminology; afterwards, we discuss the main challenges related to k-NN query computation on GPUs, as well as the spatial index and the main data structures used in this work.

3.1 Graphics Processing Units: terminology

A GPU consists of an array of n_{SM} *multithreaded streaming multiprocessor* (SMs), each with n_{core} cores, yielding a total number of $n_{SM} \cdot n_{core}$ cores. Each SM is able to run *blocks* of *threads*, namely *data-parallel tasks*, with threads in a block running concurrently on the cores of an SM. Since a block typically has many more threads than the cores available in a single SM, only subsets of threads, called *warps*, can run in parallel at a given time instant. Each warp consists of sz_{warp} *synchronous, data parallel threads*, executed by an SM according to a SIMD-like paradigm [6]. We note that in current hardware a warp is managed as a block of 32 threads.

3.2 Motivating challenges and proposal sketch

The problem of computing repeated k-NN over massive moving object observations is mainly characterized, from the GPU perspective, by the indeterminateness affecting the spatial extension of the queries, since this depends on local objects densities. The usual approach to reduce the amount of computations per query is to adopt a spatial index based on some kind of tree. To compute a k-NN query one then has to perform a tree visit, exploring only parts of the tree corresponding to regions possibly enclosing nearest neighbours. Since the spatial extension of each k-NN query is unknown, different queries possibly require to visit different paths inside the tree or different amounts of leaves. Moreover, depending on the kind of tree used, each leaf possibly contains different amounts of objects with respect to other leaves, thus strengthening the challenge of materializing uniform GPU workloads. In other words, the problem is to batch enough work per GPU SM while entailing uniform workloads to improve processor occupancy, limit branch divergence and ensure locality-preserving access to the GPU memory.

Our proposal tries to tackle these issues by exploiting an iterative approach in our GPU processing pipeline. First, we adopt a spatial index based on point-region quadtrees, where blocks of close objects and queries are spatially reordered and stored according to the quadtree leaves; the leaves are in turn structurally defined by a Morton-based coding schema which permits to easily visit the tree and entails spatial locality during the processing. Second, we iteratively process all the queries in parallel, where at each iteration the various blocks of close objects are scanned in parallel to update the result sets of corresponding blocks of k-NN queries.

3.3 Space partitioning and indexing

Grid-based spatial indexes are generally reported to have low construction and maintenance costs, making them suitable for update-intensive settings [2, 7], especially when exploiting GPUs [8]. Among these, uniform grids may fail in guaranteeing the generation of balanced workloads when processing skewed spatial distribution, since parallel tasks would be assigned to uneven blocks of objects/queries indexed by the various grid cells. Other regular grid-based indexes, such as *quadtrees*, allow achieving much better workload balance with relatively small construction/maintenance costs. Moreover, they exhibit nice mathematical/structural properties which fit quite well massively parallel architectures such as the GPUs. Indeed, we adaptively partition the Minimum Bounding Rectangle (MBR) $\mathcal{G}$, containing all the objects positions during any tick, into a set of variably sized cells belonging to a grid $\mathcal{C}$ induced by a *point-region quadtree*, whose cells correspond to quadtree leaves. Each cell $c \in \mathcal{C}$ is mapped to a unique integer ID induced by a Morton coding schema, such that the total order given by IDs on cells preserves spatial locality. Based on the spatial partitioning $\mathcal{C}$, we therefore associate each object location $p \in P$ and k-NN query $q \in Q$ with the enclosing cell $c \in \mathcal{C}$.

3.4 Relevant data structures

In this section we review the main data structures used during the query processing.

3.4.1 General overview

data structures containing objects or queries information use the *structure of vectors* (*SoV*) layout. This layout gives remarkable benefits when designing GPU algorithms, above all code reuse and efficient interplay between different operations carried on GPU [9, Ch.33]. Moreover, it facilitates the exploitation of data locality and the use of coalescing or caching, whenever possible, thus offering substantial chances to boost the overall memory throughput, which is paramount when designing GPU algorithms.

3.4.2 k-NN queries result set layout

the result set of each k-NN query (also called *nearest neighbours list*) has fixed size k and shall be stored in global memory. We arrange such queries lists linearly, as depicted in Figure 1. This layout, coupled with sequential

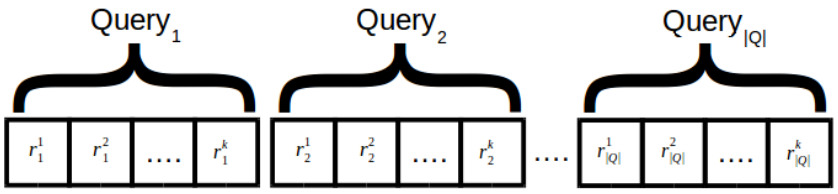

Figure 1: Result set layout. r_i^j denotes j-th result of query q_i.

accesses, proves to be quite effective in terms of exploiting caching capabilites, at least until k does not get large. In the latter case, to still benefit from the linear layout we also propose an alternative strategy, which updates the result lists in global memory by exploiting coalesced access patterns. Indeed, when k is large this technique turns out to be better, although it implies a slightly increased computational overhead. We defer the details of this technique to Section 4.2.

4 Processing Pipeline

The pipeline can be described in terms of a succession of three macro stages: *(i) index creation*, *(ii) moving objects indexing*, and *(iii) iterative query processing*. More precisely, the data entering the processing pipeline at the end of each tick are first processed to select the index parameters and build an empty index (stage *(i)*). Then objects/queries get associated with index cells and are subsequently sorted, so that those contained within the same cell are stored in contiguous memory locations (stage *(ii)*). The subsequent phase (stage *(iii)*), which essentially represents the pipeline's core, processes the k-NN queries.

When processing repeated k-NN queries the same procedure is repeated for each tick, except for stage *(i)* which is skipped when the objects spatial distribution does not change significantly with respect to recent, previous ticks. Thus, for the sake of readability, hereinafter we omit the subscript that indicates the tick, and denote by P, Q, and R, respectively, the up-to-date object positions, the queries, and the result set associated with a generic tick.

4.1 Index Creation and Moving Objects Indexing

4.1.1 Index Creation

we rely on a point-region (PR) quadtree as the backbone of the spatial index, exploiting PR-quadtrees ability to partition the space in differently sized parcels containing similar amounts of entities. Observing that - most of the times - space distributions do not change their characteristics dramatically over short time intervals, the index is rebuilt only whenever needed, i.e., when we detect that the overall amount of computations yielded during the last tick exceeds by a given factor the amount of computations yielded during past, recent ticks. The goal of this phase, consequently, is to create - almost entirely on GPU - a space partitioning $\mathcal{C}$ over $\mathcal{G}$ where each cell of $\mathcal{C}$ is a leaf of the PR-quadtree that does not contain more than th_{quad} objects. In the following we describe the PR-quadtree construction procedure.

First, we fix a maximum quadtree depth, l_{max}, and consider a *virtual full quadtree* whose leaves are all at the same level l_{max}. Incidentally we observe that, in terms of space partitioning, this is equivalent to a *uniform grid* of $2^{l_{max}} \times 2^{l_{max}}$ cells. Then, we compute in parallel the Morton codes of the objects at level l_{max} - this corresponds to computing the indices of the quadtree quadrants at level l_{max} where objects fall - (operation carried on **GPU**) and sort the objects accordingly (by means of Radix Sort, **GPU**). We perform such operations to exploit the direct relationship between Morton codes and quadtrees structural properties: given any Morton code z associated with a quadrant c at level l_{max} – where c is thus identified by the pair (l_{max}, z), we can determine the Morton code z' of a quadrant c' at any level $l \leq l_{max}$ - where c' is thus identified by the pair (l, z') and c' spatially includes c, through a simple (bitwise) operation $z' = \lfloor \frac{z}{4^{l_{max}-l}} \rfloor$. Another property is that the initial ordering of objects obtained by sorting the Morton codes at level l_{max} is invariant for any level $l \leq l_{max}$. In other words, the initial sorting suffices to guarantee that all the objects falling in a given quadrant of the tree (at any level $l \leq l_{max}$) are stored in contiguous memory locations. Therefore, after the initial sorting, without loss of generality we can identify the set of object positions P', $P' \subseteq P$, included in a given quadrant of the quadtree by using the *interval of memory indexes* $I_{P'} = [\delta, \delta + |P'|)$ where they are actually stored.

Aftwerwards, the algorithm starts iteratively to build the *actual* quadtree, level by level. Let us suppose that I_A represents the set of intervals related to quadrants which need to be split. Initially, $I_A = \{\ [0, |P|)\ \}$. Then, at each iteration, for each $(l-1)$-level quadrant added to I_A for splitting, we first identify the starting/ending positions of the four intervals of the l-level quadtree quadrants, and store these intervals in a temporary variable (operation carried on **GPU**). Then, we determine which quadrants need further splitting at next level (their intervals are added to I_A, since the corresponding quadrants contain more than th_{quad} objects) and which ones represent final leaves (their identifiers are added to $\mathcal{C}$) (operation conveniently carried on **CPU**). The process ends whenever no more quadrants need to be split, and thus I_A turns out to be empty, or the maximum possible quadtree level l_{max} is reached (in this case the remaining quadrants are added to $\mathcal{C}$).

The complexity of the index creation is approximatively linear in the amount of objects. We defer to the Appendix further details on the matter.

Building a lookup table to map coordinates to cells. The usual approach for finding a quadtree leaf would consist in traversing the tree from the root, recursively choosing relevant nodes until the target leaf is reached. Unfortunately, on GPU this approach entails repeated irregular memory accesses and a non predictable number of operations for each leaf search. For this reason we use a different approach, characterized by a slightly larger memory footprint.

Let us suppose that the deepest level created in a quadtree $\mathcal{C}$ is l_{deep}, $l_{deep} \leq l_{max}$. We virtually divide the space covered by $\mathcal{C}$ according to a uniform squared grid composed of $2^{\cdot l_{deep}} \times 2^{\cdot l_{deep}}$ cells, and denote it by $\mathcal{C}^{l_{deep}}$. We note that the space partitioning induced by $\mathcal{C}^{l_{deep}}$ also corresponds to the full quadtree having depth l_{deep}. Thanks to PR-quadtree structural properties, any quadtree leaf at a level l, $l \leq l_{deep}$, corresponds to the union of $4^{(l_{deep}-l)}$ contiguous cells of $\mathcal{C}^{l_{deep}}$. Therefore, a mapping between $\mathcal{C}^{l_{deep}}$ and $\mathcal{C}$ cells can be easily established by means of a lookup table that realizes the function $z_{map} : \mathcal{C}^{l_{deep}} \rightarrow \mathcal{C}$, which associates each cell in $\mathcal{C}^{l_{deep}}$ with the cell in $\mathcal{C}$ containing it. This allows to retrieve the quadtree leaf that contains any entity with simple numeric operations in constant time, making it a lightweight operation thanks to the structural simplicity underlying uniform grids.

At this point we can also highlight how the enumeration given to quadtree leaves actually establish a total order. This can be demonstrated by associating each leaf (l, z) with the Morton code of the *first* (decreasing with respect to the Morton order) quadrant it covers at level l_{deep}, i.e., by computing $z'' = z \cdot 4^{l-l_{deep}}$. Figure 2 shows a simple example ($l_{deep} = 2$), where the total order induced by z'' over the quadrants of $\mathcal{C}$ is illustrated in the left picture by the dashed line. In the right picture we see the same dashed line, yet going on at l_{deep}, reflecting the total order through which the identifiers of the $\mathcal{C}^{l_{deep}}$'s cells are arranged in z_{map}.

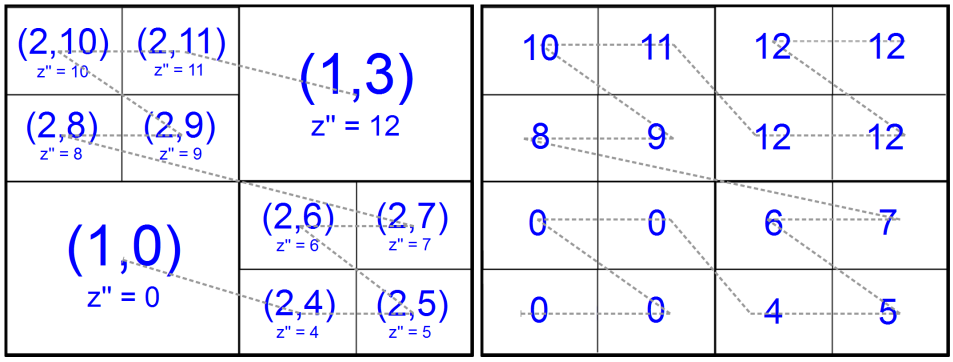

Figure 2: Example of the mapping established by z_{map} between the quadtree-induced grid $\mathcal{C}$ (left side) and the uniform grid $\mathcal{C}^{l_{deep}}$ (right side) related to the quadtree deepest level.

We have to mention that the performance related to lookups in z_{map} heavily depends on the ability to exploit the GPUs caching capabilities. Indeed, z_{map} may have a relevant size - depending on l_{deep}. In light of this, it is important to design the *memory layout* of z_{map} to enhance data locality - this is defined by the Morton order at l_{deep} - as well as to sort objects and queries according to the same order structurally defining z_{map} before the actual lookups happen (see Section 4.1.2). The cost of initializing z_{map} (done on **GPU**) after the index creation is, in practical terms, negligible.

4.1.2 Moving objects indexing and active cells determination

during this phase we assign each moving object to a specific quadtree leaf. More precisely, first each object in P is associated with a cell in $\mathcal{C}^{l_{deep}}$. Subsequently, objects are *sorted* according to the identifiers of the cells to exploit caching when subsequently accessing z_{map}. Finally, z_{map} is used to retrieve the identifier (l, z) of the corresponding cell in the quadtree-induced $\mathcal{C}$. Finally, the resulting, sorted, struct of vectors holding P is *indexed* to determine the positions of the first and last object belonging to cells of $\mathcal{C}$ that enclose at least one object. We define these cells as *active*, and denote the set of active cells by $\hat{\mathcal{C}}$, where $\hat{\mathcal{C}} \subseteq \mathcal{C}$.

The complexity of this subphase is linear with respect to the amount of objects. We defer further details to the Appendix.

4.2 Iterative k-NN queries computation

Once the set of active cells $\hat{\mathcal{C}}$ is determined, the actual query computation may start. Since we rely on a spatial index based on a PR-quadtree, for each query we have to perform a *tree visit* to compute the list of k nearest neighbours. To exploit the computational power of GPUs we need to create GPU workloads such that computation is distributed uniformly across GPU SMs and objects, queries and results data are arranged so that *coalescing* and *caching* are exploited as much as possible. Further, since the tasks to perform are dependent on data, we have to create those workload on-the-fly, as tree visits progress.

The underlying idea is to achieve these goals by means of an *iterative* approach, exploiting spatial proximity between nearby queries as visits progress. We distinguish between the operations carried on during the *first iteration* and those carried on during *subsequent iterations*.

4.2.1 First iteration

the first iteration orchestrates the work associated with the beginning of the queries tree visits. This corresponds to associating each query with the cell (quadtree leaf) in which its center falls, and subsequently computing the

distances between the query center and the objects enclosed by the cell. The first iteration ends by updating the queries nearest neighbours lists according to the distances computed. We structure the first iteration in two smaller phases: (i) *query indexing and GPU task materialization* and (ii) *distance computations* between queries and objects.

Query indexing and task materialization. Since each k-NN query is represented by a point, queries are mapped to specific grid cells. Consequently, such phase is equivalent to the moving objects indexing phase (Section 4.1.2). Since it is meaningless to consider devoid cells for the purposes of distance computations, we consider only the set of non-devoid cells containing at least a query. We denote by $\overline{\mathcal{C}}$ such set, noting that $\overline{\mathcal{C}} \subseteq \hat{\mathcal{C}} \subseteq \mathcal{C}$. Each cell in $\overline{\mathcal{C}}$ will thus be assigned to a multithreaded task, to be scheduled to a distinct GPU SMs, where each concurrent thread is in charge of a distinct query.

Prior to the *distance computations* phase, we finally sort in descending order the set of tasks on the basis of their *computational weight*, i.e., the amount of distances to be computed: by scheduling tasks according to this order we reduce possible workload imbalance across the GPU SMs.

Distance computations. At this point the goal is to compute, for each query, the list of (up to) k nearest objects within the assigned cell and update the lists accordingly. We saw above that the end product of the query indexing and sorting operations consists in a set of tasks, one per active cell enclosing at least one object, representing the GPU workload in charge of computing such lists. This allows to conveniently pack together computations related to spatially nearby entities, allowing to reduce the overall amount of computations and to exploit data locality.

Accordingly, our approach relies on two pillars. The first one is represented by a GPU-friendly *k-selection* algorithm based on buckets [10]. Starting from a set of objects, this algorithm allows to find the k nearest objects without having to explicitly store and sort distances in memory, thus reducing the overall complexity in terms of time and space. The second pillar is represented by a proper access pattern which allows to maximize the memory throughput when updating the queries nearest neighbours lists. Considering the linear layout used for the queries result set (Figure 1), different approaches may represent the best choice depending on k.

Distance computations – "*cached writes*" approach. In this paragraph we focus on a strategy which relies only on GPU caching capabilities when updating the queries nearest neighbours lists. Algorithm 1 reports such strategy, where **local** keyword denotes automatic variables stored in private registers.

Each parallel task in charge of a given $c \in \overline{\mathcal{C}}$ is assigned to a specific SM (line 2) and executed according to a per-query parallelization (line 3). Each thread first loads query information (line 3), i.e., query coordinates and the associated cell identifier. Since queries belonging to the same cell are stored in contiguous memory locations, accesses to query data are *coalesced* during the first iteration. Subsequently, each thread finds out the minimum and maximum distance ($dist_{min}$ and $dist_{max}$ respectively) between the query center and the objects within the cell (function *findMinMaxDist*, line 4). This is achieved through a simple scan over the set of objects enclosed by c. Since every thread in a warp perform such scan by accessing objects data in the same order, and considering that different warps in a block access nearby objects, this behaviour exhibits strong locality of reference, and is thus suitable to fully exploit the GPU *caching* capabilities. Once $dist_{min}$ and $dist_{max}$ are determined, the algorithm goes on by finding a distance which encompasses only the k nearest neighbours within c, $dist_k$. This is achieved by calling the *findKDist* function (line 5), which implements a k-selection algorithm based on buckets [10], iteratively going on until a suitable distance is found. We note that whenever the amount of objects in c is less than k, the function can immediately return $dist_k = +\infty$ without performing any computation. Even inside *findKDist*, threads of the same warp access c's objects in the same order; since different warps access objects arranged in nearby memory locations, this entails again an efficient usage of GPU caches.

Once $dist_k$ is determined, each thread can actually start writing sequentially the list of nearest neighbours associated with the query, according to the layout shown in Figure 1 (line 7). The thread terminates once it writes out the distance of the farthest object in the list, $MAXDIST_q$, and the actual amount of results written in the list, $NUMRES_q$ (lines 12 and 13), both stored in global memory. We note that such writes are coalesced, thanks to per-query parallelization.

Distance computations – "*coalesced writes*" approach. This approach is equivalent to the one presented above except for the access pattern used to update queries lists. More precisely, once we have computed $dist_k$ for each query, the key idea is to parallelize at warp level by assigning each query to a *warp*. In turn, inside each warp we parallelize distance checks at thread level, while writes related to lists updates are cooperatively orchestrated at warp level, thus allowing to exploit coalescing when flushing out data in global memory. It is evident that such cooperative strategy requires, at *task level*, the usage of a temporary buffer stored in shared memory to accumulate

Algorithm 1: $distComp(\overline{\mathcal{C}}, Q, P, k)$

Input :

- The set of active cells with at least one query, $\overline{\mathcal{C}}$.
- The reordered query set *Q* and object set *P*, along with the indexing information associated after the respective sorting phases.
- The size of the queries neighbours lists, *k*.

Output:

- The struct of vectors containing the query result set, $(ID, DIST)$.
- The vector containing the maximum distance detected for each query, *MAXDIST*.
- The vector containing the amount of nearest neighbours found for each query, *NUMRES*.

```
1  begin
2      foreach c ∈ C̄ parallel_block do
3          foreach q ∈ c parallel_thread do
4              local (dist_min, dist_max) ← findMinMaxDist(q, c)
5              local dist_k ← findKDist(q, c, dist_min, dist_max, k)
6              local i ← 0, maxdist ← 0
7              foreach p ∈ c do
8                  if ((dist(p, q) < dist_k) ∧ (q_id ≠ p_id)) then
9                      maxdist ← max(maxdist, dist(p, q))
10                     (ID_q[i], DIST_q[i]) ← (p_id, dist(p, q))
11                     i ← i + 1
12             MAXDIST_q = maxdist
13             NUMRES_q = i
```

results, coupled with proper management operations. To this end, we exploit the native CUDA *ballot* function which allows to find out (i) who found a nearest neighbour across a warp (if any), and (ii) manage writes inside the shared memory buffer without having to recur to synchronizing barriers. This approach trades a slightly increased computational complexity for a substantial increase in memory throughput when *k* gets large.

Distance computations – Complexity. The complexity of the distance computation phase is mainly dictated by the k-selection algorithm, whose number of iterations depends on the spatial distribution inside a cell. We defer to the Appendix further details.

4.2.2 Subsequent iterations

after the first iteration each query is associated with a list of nearest neighbours, containing up to *k* objects contained within the same cell where the query center falls. Depending on the objects spatial distribution and the derived index, however, it is very probable that a substantial fraction of queries have an incorrect or incomplete list. This happens when the result of a query contains at least one object falling outside the cell of the query, and such objects are nearer than the farthest object in the list computed so far. Another trivial case is when a query falls inside a cell containing less than *k* objects (and $k < |P|$). Clearly, we need to look at other neighbouring cells to complete the list. From now on, we call such queries *active queries*. For each one of these we need to perform a tree visit, starting from the leaf in which their center fall, considering those leaves whose spatial extent may contain potential nearest neighbours, and update the query list accordingly as the visit progresses. Considering the potential amount of active queries to be processed after the first iteration, the main challenge is to devise a strategy able to batch the work resulting from such tree visits to generate workloads suitable for GPU processing.

Our proposal is based on an iterative approach, where tree visits and distance computations related to spatially nearby active queries are packed together, iteration after iteration, until no query remains active, i.e., every query is associated with the final correct list of nearest neighbours.

At query level, the basic idea is to perform the aforementioned tree visit, starting from the leaf $c \in \mathcal{C}$ in

which the query center falls, by splitting it into two sub-visits. Since quadtree leaves are totally ordered due to the Morton-induced encoding, as illustrated in Figure 2 (left picture), one sub-visit proceeds towards "left" (with respect to the order) while the other one proceeds towards "right", until a quadtree leaf that might contain potential nearest neighbours can be found, i.e., there is at least one non-empty leaf whose borders are *nearer* than the *farthest* object in the query list computed so far. The idea is to exploit the spatial locality-preserving property of Morton codes to visit, first, those leaves which likely happen to be closer to the query center. In this way, we can minimize the amount of cells to be inspected for potential objects to be returned by a query, pruning most of them on the basis of the current $dist_k$ and the borders of the inspected cells. To further maximize this chance of pruning, the two subvisits are conducted in an alternate fashion, as long as a query remains active in both directions.

In the following we sketch out our strategy in more detail.

1 – Iterative processing, massive tree navigation. At the beginning of each iteration, for each query we alternatively consider one navigation direction, either left or right, to visit the quadtree leaves as long as possible nearest neighbours may be found in at least one direction, and update the navigation status of each active query accordingly. At the end of this operation, each active query is assigned to the first unvisited leaf, in the direction considered, whose spatial extent may contain possible nearest neighbours. If no useful leaf is found along a direction, the query is flagged as *inactive* with respect to that direction.

As regards the GPU implementation, the main challenge is to conveniently perform such operation by assigning each active query to a single GPU thread. Indeed, if we end up processing spatially faraway queries in the same SM, we may encounter two serious performance issues: the first one is due to *execution branching* inside warps (substantial different quadtree visits may entail different flows of operation execution), while the second originates from memory issues (the quadtree information accessed by the SM threads may be stored in faraway memory locations, thus denying possible caching benefits). To this end, we exploit the fact that active queries come already sorted from previous iterations (see Section 4.2.1 and the query sorting and task materialization operation described below), i.e., we assign active queries to threads according to the order achieved by means of sorting during the previous iteration and characterized by the same visit direction. Consequently, the aforementioned issues are implicitly addressed, since spatially nearby queries likely happen to be arranged in nearby memory locations due to Morton codes properties.

At thread level, we have to update the navigation status of a query restarting from the *previously assigned* leaf. To this end, we consider again the virtual full-quadtree having depth l_{deep}, $\mathcal{C}_{l_{deep}}$, instead of the actual one. In this sense, we (possibly) assign a quadtree $\mathcal{C}$ leaf to a query only whenever we detect a quadrant at l_{deep} whose borders are closer than the farthest object in the query list (at distance $dist_k$ from the query). If this happens, we then just need to perform a single lookup in z_{map}, similarly to what is done in the moving objects and query indexing phases, to retrieve the enclosing leaf of $\mathcal{C}$. This design choice gives a big major advantage from the GPU perspective, in that the virtual full-quadtree, thanks to its structural regularity, can be navigated without performing any lookup in memory, since the spatial extension of its quadrants can be determined mathematically on-the-fly by each thread. The only lookup needed is the one we have to perform when we detect a quadrant of interest in l_{deep}. Secondarily, since queries are assigned to threads according to the sorting order yielded during the previous iteration characterized by the same visit direction, we exploit data locality when accessing z_{map}.

Query-wise, the complexity of this subphase is dictated by the amount of active queries in the considered direction. At query level, the worst case scenario corresponds to visiting all $\mathcal{C}_{l_{deep}}$ nodes.

2 – Iterative processing, query sorting and task materialization. Subsequently, queries are *sorted* according to the newly assigned cell identifiers. The same sorting operation also partitions between active and inactive queries as well: once active queries get sorted it suffices to find the first query having the inactive flag set to determine the extent of both sets. Such simple, yet massively parallel, operation is conveniently performed on GPU.

After the sorting operation we only focus on those queries still considered active. By virtue of sorting, the queries associated with the same cell are displaced in contiguous memory locations. We exploit again this property, as in Section 4.2.1, to determine the first and last query for each cell having at least one query assigned. The outcome is represented again by a set of tasks, one per active cell associated with at least one active query. We denote such set by $\overline{\mathcal{C}}$, as done previously.

3 – Iterative processing, nearest neighbours lists update. Afterwards, we proceed by updating the result lists of all the active queries by considering their newly assigned index cells. This step is almost equivalent to the one described in Algorithm 1: for each query, first we find out the new $dist_k$, by considering the objects contained in the current result list of the query and those occurring in the newly assigned cell (this is done again by means of the k-selection algorithm). Finally, we update the list according to the new $dist_k$. Here it is worth remarking that for each query we can *prune* out consistent amounts of objects in the newly assigned cell simply by looking

at the distance of the farthest object in the result list found so far. Obviously, such optimization cannot be used in case a query has less than k neighbours in the list. The complexity of this subphase is similar to the **Distance computation** one, performed during the first iteration (Section 4.2.1). We defer further details to the Appendix.

5 Experimental Setup

All the experiments are conducted on a PC equipped with an Intel Core i7 2600 CPU, running at 3,4 GHz, with 16 GB RAM and a Nvidia GTX 580 GPU with 3 GB of RAM coupled with CUDA 5.5. The OS is Ubuntu 12.04. We exploit a publicly available framework [2] for both workload generation and testing. We use three types of synthetic datasets: *uniform datasets*, in which moving objects are distributed uniformly in the space, *gaussian datasets*, in which moving objects gather around *hotspots* by following a normal distribution, and *road network datasets*, where objects are distributed uniformly over the edges of a network (in our case, the San Francisco road network) and move along the edges. In the gaussian datasets case the skewness depends on the amount of hotspots: the more they are, the more the resulting distribution is uniform. In all tests we compute repeated k-NN queries over 30 ticks; each query is issued by a moving object - according to the problem statement introduced in Section 2.3. To model object movements the framework generates 30 instances of each dataset, one for each tick. The overall amount of queries per tick is equal to the amount of moving objects (i.e., one query per object). Table 1 summarizes the main parameters used to generate the datasets. The framework uses a generic spatial distance unit u (e.g., meters). The usage of synthetic (or partly synthetic) datasets only is consistent with previous literature [2] and can be motivated by observing - apart from the lack of suitable real-world datasets - that the datasets characteristics space possibly influencing the algorithms performance is very vast; as a consequence, exploring effectively such space is feasible only if one has full control over such factors.

We use K-NN$_{\text{GPU}}$ to denote our proposal. We denote by K-NN$^{CACHE}_{\text{GPU}}$ the flavour using the *cached writes* approach, while K-NN$^{COALESCE}_{\text{GPU}}$ uses the *coalesced writes* approach. Wherever not specified, K-NN$_{\text{GPU}}$ = K-NN$^{CACHE}_{\text{GPU}}$. As CPU sequential competitor we consider the sequential k-NN search algorithm offered by the well-established FLANN library [1] and is denoted by K-NN$_{\text{CPU}}$. The GPU-based baseline is the brute-force algorithm presented in [4] and is denoted by K-NN$_{\text{BASELINE}}$.

Spatial region	All tests occur in a squared spatial region with side length of 22500 u.
Objects maximum speed	In all tests the maximum speed of each object is fixed to 200 u per tick (Δt), where the objects are allowed to change their speed as described in [2]. In general, changes in speed may slightly alter the objects distribution but do not change the distribution general properties.
Query rate	The percentage of objects that issue a k-NN query during every tick is always set to 100%.

Table 1: Relevant data and workload generation parameters.

6 Experimental Evaluation

The experimental studies conducted for this work are the following:

- *S1*: We study how K-NN$_{\text{GPU}}$'s performance is affected when varying the maximum amount of objects per quadtree leaf (thus influencing the tree height), the amount of results per query (k), and the dataset skewness.
- *S2*: We compare K-NN$_{\text{GPU}}$ against K-NN$_{\text{BASELINE}}$.
- *S3*: We compare K-NN$_{\text{GPU}}$ against K-NN$_{\text{CPU}}$.

S1 – Tree height, neighbours list size k and spatial skewness impacts on K-NN$_{\text{GPU}}$'s performance

Tree height and neighbours list size. When processing any dataset, two crucial parameters are represented by the tree height, which is controlled indirectly by altering the maximum amount of objects admitted in a single

[1] http://www.cs.ubc.ca/research/flann/, version 1.8.4.

quadtree leaf (th_{quad}), and the neighbours list size k associated with the each query. Choosing an optimal th_{quad} is strictly connected to k: if the tree height is too high with respect to k, then many leaves with few objects shall be visited, thus increasing the amount of operations needed to perform such operations. On the other hand, if the tree height is too low we end up visiting few leaves with many objects, thus possibly performing a relevant amount of useless computations due to the reduced pruning power of the index. As a consequence, it is necessary to find an appropriate th_{quad} for a given k.

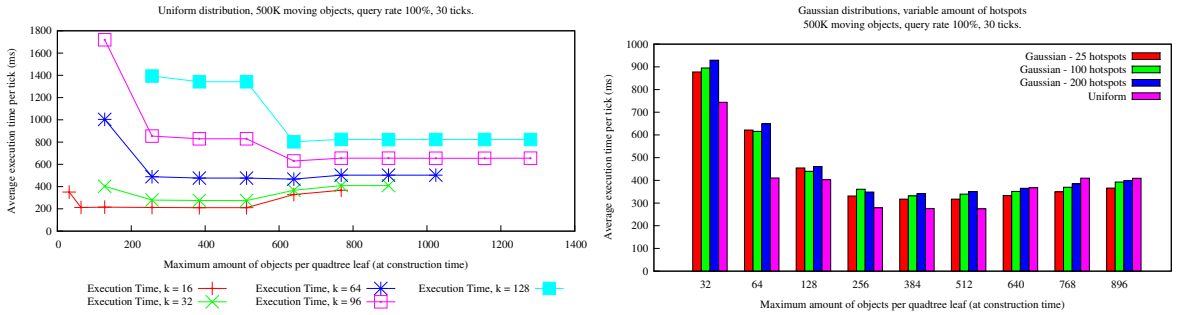
Uniform distribution, 500K moving objects, query rate 100%, 30 ticks.
Average execution time per tick (ms)
Maximum amount of objects per quadtree leaf (at construction time)
Execution Time, k = 16
Execution Time, k = 32
Execution Time, k = 64
Execution Time, k = 96
Execution Time, k = 128
Gaussian distributions, variable amount of hotspots
500K moving objects, query rate 100%, 30 ticks.
Gaussian - 25 hotspots
Gaussian - 100 hotspots
Gaussian - 200 hotspots
Uniform

Figure 3: *Left plot*: Relationship between tree height (indirectly controlled) and k, and its repercussions on K-NN$_{\text{GPU}}$'s performance. *Right plot*: Skewness repercussions on execution time and th_{quad}'s optimality.

Figure 3, *left plot*, shows how each k is associated with a range of optimal th_{quad} values for which the algorithm's execution time is minimized. We observe how such ranges are relatively wide as well, a property which is desirable since this minimizes performance fluctuations even when not using an optimal th_{quad}. Finally, even if it is not evident for every curve the execution time starts increasing again whenever th_{quad} gets too high with respect to k, since the pruning power of the index gets reduced. Incidentally, we observe how the execution time increases whenever k increase. We note that the same performance trends can be replicated with skewed distributions (such as gaussian and network ones), even though characterized by slightly higher execution times.

Spatial skewness impact on finding an optimal th_{quad}. In this batch of experiments we want to check whether the skewness has relevant impacts on finding an optimal th_{quad} range. Datasets are distributed according to a gaussian distribution, each characterized by a different amount of hotspots to yield differently skewed distributions. All datasets have a fixed amount of objects equal to 500K; other dataset characteristics are set to their defaults, according to Table 1. The size of nearest neighbours lists is set to $k = 32$. From Figure 3, *right plot*, we see how the skewness does not influence, if not marginally, the optimal th_{quad} range for a given k. Even if in this series of experiments we use a fixed k, it is possible to show that the performance trends associated with different k values are approximately the same.

S2 – K-NN$_{\text{GPU}}$ vs K-NN$_{\text{BASELINE}}$

In this study we take into consideration two main parameters, i.e., the *amount of objects* and the amount of nearest neighbours per query, k. In light of the results shown in study S1, for what regards K-NN$_{\text{GPU}}$ we set $th_{quad} = 12k$ when $32 \leq k \leq 256$ since this assures, on average, the best performance, while we use $th_{quad} = 192, k < 32$ and $th_{quad} = 2048, k > 128$.

In the first batch of experiments we consider a uniform distribution, where we vary the amount of moving objects and set $k = 32$, while other dataset characteristics are set according to defaults. From Figure 4, *left plot*, we see how K-NN$_{\text{GPU}}$ heavily outperforms K-NN$_{\text{BASELINE}}$ as soon as the amount of objects gets relevant. In the second batch of experiments we still use a uniform distribution and vary the amount of nearest neighbours per query (k) while keeping fixed the amount of moving objects (100K). Other characteristics are set to defaults. From Figure 4, *right plot*, we see how K-NN$_{\text{BASELINE}}$'s execution time is fixed, since it depends exclusively on the amount of distances to compute and sort, while K-NN$_{\text{GPU}}$'s execution time increases when k increases - as expected - thus reducing its advantage gap.

S3 – K-NN$_{\text{GPU}}$ vs K-NN$_{\text{CPU}}$

In this section we compare K-NN$_{\text{GPU}}$ against K-NN$_{\text{CPU}}$. K-NN$_{\text{CPU}}$ relies on kd-trees for computing k-NN queries. In our tests we force K-NN$_{\text{CPU}}$ to use 1 CPU core and an optimized L2 distance functor. We also set to 32 the

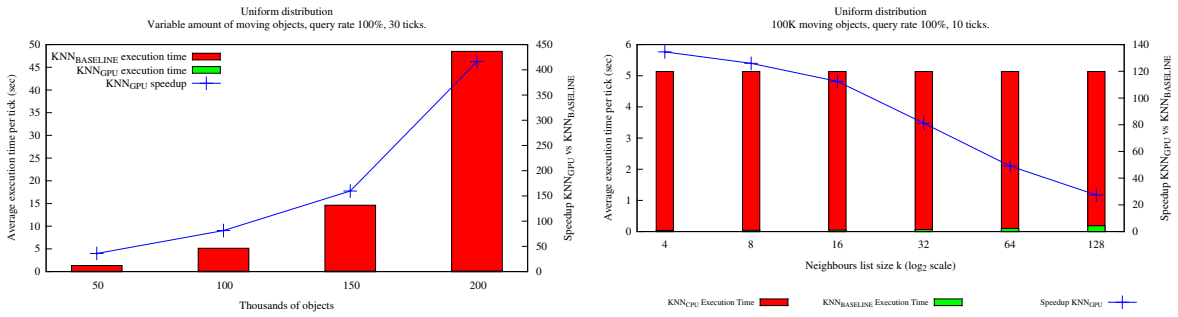
Uniform distribution
Variable amount of moving objects, query rate 100%, 30 ticks.
KNN_BASELINE execution time
KNN_GPU execution time
KNN_GPU speedup
Average execution time per tick (sec)
Speedup KNN_GPU vs KNN_BASELINE
Thousands of objects
Uniform distribution
100K moving objects, query rate 100%, 10 ticks.
Average execution time per tick (sec)
Speedup KNN_GPU vs KNN_BASELINE
Neighbours list size k (log2 scale)
KNN_CPU Execution Time
KNN_BASELINE Execution Time
Speedup KNN_GPU

Figure 4: K-NN$_{\text{GPU}}$ vs K-NN$_{\text{BASELINE}}$. *Left plot*: variable amount of moving objects, $k = 32$. *Right plot*: variable nearest neighbours list size k, fixed amount of objects (100K).

maximum amount of objects per kd-tree leaf, since this value proves to be the best choice in our experimental setting. For what regards K-NN$_{\text{GPU}}$, we use the same settings used in S2 for th_{quad}. In this analysis we consider datasets characterized by different spatial distributions and different amounts of moving objects. Also, we study how K-NN$_{\text{GPU}}$ behaves according to different amounts of nearest neighbours per query (k).

Varying the amount of moving objects. In the following batch of experiments we study K-NN$_{\text{GPU}}$ and K-NN$_{\text{CPU}}$ when varying the amount of moving objects. The dataset characteristics varied across the experiments are (i) the spatial distribution and (ii) the amount of moving objects (between 100K and 1500K). In the Gaussian case the amount of hotspots is fixed to 25, thus yielding a moderately skewed distribution, while other characteristics are set to defaults.

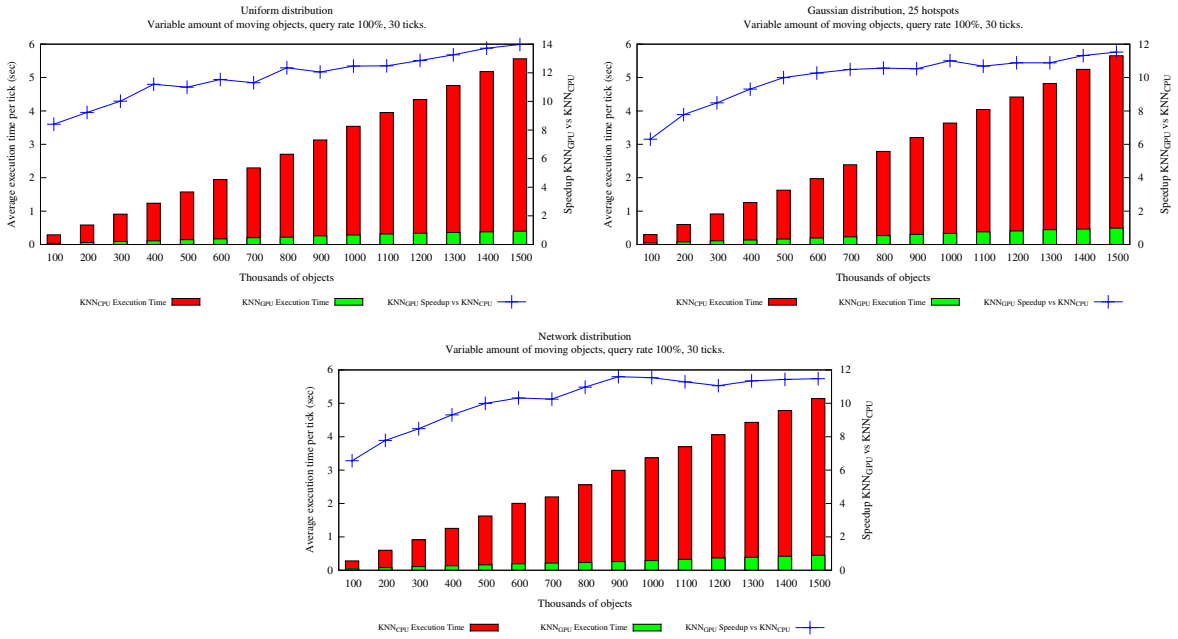
Uniform distribution
Variable amount of moving objects, query rate 100%, 30 ticks.
Average execution time per tick (sec)
Speedup KNN_GPU vs KNN_CPU
Thousands of objects
KNN_CPU Execution Time
KNN_GPU Execution Time
KNN_GPU Speedup vs KNN_CPU
Gaussian distribution, 25 hotspots
Variable amount of moving objects, query rate 100%, 30 ticks.
Average execution time per tick (sec)
Speedup KNN_GPU vs KNN_CPU
Thousands of objects
KNN_CPU Execution Time
KNN_GPU Execution Time
KNN_GPU Speedup vs KNN_CPU
Network distribution
Variable amount of moving objects, query rate 100%, 30 ticks.
Average execution time per tick (sec)
Speedup KNN_GPU vs KNN_CPU
Thousands of objects
KNN_CPU Execution Time
KNN_GPU Execution Time
KNN_GPU Speedup vs KNN_CPU

Figure 5: K-NN$_{\text{GPU}}$vs K-NN$_{\text{CPU}}$, variable amount of objects, fixed $k = 32$.

From Figure 5 we see how, in general, K-NN$_{\text{GPU}}$'s speedup increases whenever the amount of objects increases, since the amount of calculations (in terms of distances to compute) increases, thus making the query processing more and more a compute-intensive task - therefore favouring K-NN$_{\text{GPU}}$. The other main observation relates to the skewness: speedups achieved with skewed distributions are slightly lower than those achieved with uniform distributions; this fact is expected, since the indexing used by K-NN$_{\text{GPU}}$ cannot totally avoid imbalances between single tasks.

Varying the neighbours lists size, *k*. In this batch of experiments we study how K-NN$_{\text{GPU}}$compares with respect to K-NN$_{\text{CPU}}$ when varying the neighbours list size k and the spatial distribution. This time we use a fixed amount

of objects, 1 million, while varying k in the [1,512] range. All other dataset characteristics are set to defaults.

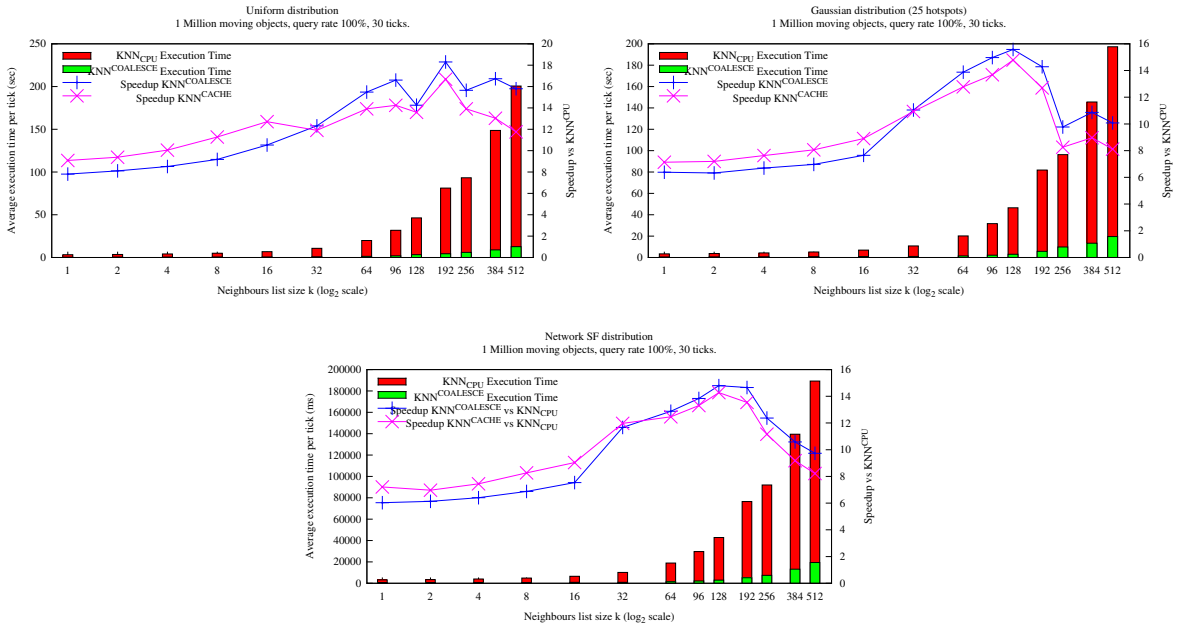
Uniform distribution
1 Million moving objects, query rate 100%, 30 ticks.
Gaussian distribution (25 hotspots)
1 Million moving objects, query rate 100%, 30 ticks.
Network SF distribution
1 Million moving objects, query rate 100%, 30 ticks.
KNN_CPU Execution Time
KNN^COALESCE Execution Time
Speedup KNN^COALESCE
Speedup KNN^CACHE
Speedup KNN^COALESCE vs KNN_CPU
Speedup KNN^CACHE vs KNN_CPU
Average execution time per tick (sec)
Average execution time per tick (ms)
Speedup vs KNN^CPU
Neighbours list size k (log2 scale)

Figure 6: K-NN$^{COALESCE}_{\text{GPU}}$ vs K-NN$^{CACHE}_{\text{GPU}}$ vs K-NN$_{\text{CPU}}$, variable k.

From Figure 6 (in the plot the default variant used is denoted by K-NN$^{CACHE}_{\text{GPU}}$) we see how increasing k up to a certain point has positive effects on the performances, since the processing becomes more and more compute-intensive while GPU caching is still effective. However, when k gets very high we observe performance degradation due to a decreased GPU caching efficiency, which is expected considering the linear layout (Figure 1) used for the queries result set. For what relates to the skewness we observe again how it mildly influences negatively K-NN$_{\text{GPU}}$'s performances with respect to uniform distributions.

Improving the memory throughput when k gets high: from Figure 6 we observe how K-NN$^{COALESCE}_{\text{GPU}}$ outperforms K-NN$^{CACHE}_{\text{GPU}}$ whenever k is equal or greater than the size of a warp ($k \geq 32$), while we have the opposite whenever k is lower. These results can be explained by observing some key facts. First, K-NN$^{COALESCE}_{\text{GPU}}$'s strategy becomes effective once threads inside a warp get fully utilized and the required amount of nearest neighbours per query gets more and more relevant, thus making K-NN$^{CACHE}_{\text{GPU}}$'s strategy less effective in terms of memory throughput. Indeed, this combination of factors is reflected in the performance trend observed in all the figures. Second, to exploit coalescing we need a proper access pattern, which in turn requires some computational overhead to orchestrate the computations accordingly. Whenever k is low such overhead, coupled with a slight thread under-utilization when updating the queries lists, slightly penalizes K-NN$^{COALESCE}_{\text{GPU}}$.

7 Related work

In the following we review relevant works concerning low dimensional spaces, since this work considers $\mathbb{R}^2$. For what is related to traditional architectures, the vast majority of the approaches are based on the usage of kd-trees, as shown in extensive surveys such as [11, Ch. 63] or [12, Sec. 5], while other solutions are based on R-Trees [13].

To date, the first work tackling the problem of computing k-NN queries by means of a hybrid CPU/GPU approach is [4], where the authors propose a brute-force, quadratic, approach. This approach is quite simple and straightforward, yet it fits quite well the GPUs architectures and proves to be quite effective with small/medium sized datasets having high dimensionality. An improved brute-force strategy is presented in [14], where the authors use a variant of the bitonic sort algorithm, which keeps into account k when sorting the elements, to significantly reduce the overall execution time. Focusing specifically on works related to low dimensional spaces, many actually tackle the problem in static scenarios and when the problem is part of a bigger problem. Indeed, computing k-NN queries is strikingly recurrent in many fields, such as computer graphics, physics, astronomy, etc. (e.g., [15–18]). Such works almost always rely on a kd-tree based index, whose construction usually happens

on CPU. The index is subsequently used to perform the k-NN search, where the tree has to be navigated for each query: such navigation is usually performed on CPU, depending on the specific approach, while distance computations always happen on GPU. All these solutions typically exhibit a bottleneck, in that at least one core operation is performed on CPU whereas parallelizing on GPU would speed up the entire processing remarkably; moreover, these works are tailored for very specific scenarios. Other few works tackle problems more similar to the one presented here. For instance, [19] proposes an approach which uses the GPU to speed up the computation of k-NN queries over a set of static points dislocated on a given road network. The approach uses a spatial index derived from Voronoi diagrams, which turns out to be computationally expensive to compute and lookup on GPU and is therefore impractical when tackling massive amounts of points or queries. Another work [20] considers a (possibly massive) set of moving objects dislocated along a road network and, equivalently to our work, it tackles the problem of computing k-NN queries issued by a fraction of the same objects. Since the computation of the distances depends on the topology of the network, the main idea is to use the network as the underlying spatial index to limit the overall amount of computations per k-NN query. The experimental part of this work seems to indicate that the proposed approach is effective only when processing very moderate amounts of queries.

A recent work [21] targets spaces having low to moderate dimensionality (i.e., $\mathbb{R}^{4 \leq d \leq 20}$). This work is quite interesting in that, starting from a spatial index based on kd-trees, it exploits queue-based buffers to accumulate and distribute on the fly fairly uniform workloads across GPU streaming multiprocessors. Analogously to our proposal, these workloads derive from tree visits bounded to k-NN queries.

8 Conclusions

In this paper we presented a novel hybrid CPU/GPU query processing pipeline, relying on scalable grid-based spatial indexes and on an iterative approach - coupled with ad-hoc data structures and proper memory access patterns, capable of computing massive amounts of repeated k-NN queries over massive moving objects observations. The solution proposed is the first known to effectively exploit the GPUs architectural features to speed-up the query processing in such scenarios and, at the same time, contend effectively with skewed spatial distributions of objects and queries. We extensively tested our solution to study its sensitivity to parameters and data distribution. In these experiments we also prove that our solution outperforms a baseline GPU approach and achieves significant performance gains over a state-of-the-art CPU sequential competitor.

References

[1] Nikos Pelekis and Yannis Theodoridis. *Mobility Data Management and Exploration*. Springer, 2014.

[2] Benjamin Sowell, Marcos Vaz Salles, Tuan Cao, Alan Demers, and Johannes Gehrke. An experimental analysis of iterated spatial joins in main memory. *Proc. VLDB Endow.*, 6(14):1882–1893, September 2013.

[3] Victor W Lee, Changkyu Kim, Jatin Chhugani, Michael Deisher, Daehyun Kim, Anthony D Nguyen, Nadathur Satish, Mikhail Smelyanskiy, Srinivas Chennupaty, Per Hammarlund, et al. Debunking the 100x gpu vs. cpu myth: an evaluation of throughput computing on cpu and gpu. In *ACM SIGARCH Computer Architecture News*, volume 38, pages 451–460. ACM, 2010.

[4] Vincent Garcia, Eric Debreuve, and Michel Barlaud. Fast k nearest neighbor search using gpu. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*, pages 1–6. IEEE, 2008.

[5] Jim Gray and Andreas Reuter. *Transaction processing*. Morgan Kaufmann Publishers, 1993.

[6] Sunpyo Hong and Hyesoon Kim. An analytical model for a gpu architecture with memory-level and thread-level parallelism awareness. In *ACM SIGARCH Computer Architecture News*, volume 37, pages 152–163. ACM, 2009.

[7] Darius Šidlauskas, Kenneth A Ross, Christian S Jensen, and Simonas Šaltenis. Thread-level parallel indexing of update intensive moving-object workloads. In *Advances in Spatial and Temporal Databases*, pages 186–204. Springer, 2011.

[8] Claudio Silvestri, Francesco Lettich, Salvatore Orlando, and Christian S Jensen. Gpu-based computing of repeated range queries over moving objects. In *Parallel, Distributed and Network-Based Processing (PDP), 2014 22nd Euromicro International Conference on*, pages 640–647. IEEE, 2014.

[9] Matt Pharr and Randima Fernando. *Gpu gems 2: programming techniques for high-performance graphics and general-purpose computation*. Addison-Wesley Professional, 2005.

[10] Tolu Alabi, Jeffrey D Blanchard, Bradley Gordon, and Russel Steinbach. Fast k-selection algorithms for graphics processing units. *Journal of Experimental Algorithmics (JEA)*, 17:4–2, 2012.

[11] Dinesh P. Mehta and Sartaj Sahni. *Handbook Of Data Structures And Applications (Chapman & Hall/Crc Computer and Information Science Series.)*. Chapman & Hall/CRC, 2004.

[12] Edgar Chávez, Gonzalo Navarro, Ricardo Baeza-Yates, and José Luis Marroquín. Searching in metric spaces. *ACM Comput. Surv.*, 33(3):273–321, September 2001.

[13] Gísli R. Hjaltason and Hanan Samet. Distance browsing in spatial databases. *ACM Trans. Database Syst.*, 24(2):265–318, June 1999.

[14] N. Sismanis, N. Pitsianis, and Xiaobai Sun. Parallel search of k-nearest neighbors with synchronous operations. In *High Performance Extreme Computing (HPEC), 2012 IEEE Conference on*, pages 1–6, Sept 2012.

[15] Justin Heinermann, Oliver Kramer, Kai Lars Polsterer, and Fabian Gieseke. On gpu-based nearest neighbor queries for large-scale photometric catalogs in astronomy. In *KI 2013: Advances in Artificial Intelligence*, pages 86–97. Springer, 2013.

[16] Naohito Nakasato. Implementation of a parallel tree method on a gpu. *Journal of Computational Science*, 3(3):132–141, 2012.

[17] Deyuan Qiu, Stefan May, and Andreas Nüchter. Gpu-accelerated nearest neighbor search for 3d registration. In *Computer Vision Systems*, pages 194–203. Springer, 2009.

[18] Stefan Popov, Johannes Günther, Hans-Peter Seidel, and Philipp Slusallek. Stackless kd-tree traversal for high performance gpu ray tracing. In *Computer Graphics Forum*, volume 26, pages 415–424. Wiley Online Library, 2007.

[19] Marta Fort and J Antoni Sellares. Gpu-based computation of distance functions on road networks with applications. In *Proceedings of the 2009 ACM symposium on Applied Computing*, pages 1320–1324. ACM, 2009.

[20] Liao Wei, Zhang Zhiming, Yuan Zhimin, Fu Wei, and Wu Xiaoping. Parallel continuous k-nearest neighbor computing in location based spatial networks on gpus. In *Computational and Information Sciences (ICCIS), 2013 Fifth International Conference on*, pages 271–274. IEEE, 2013.

[21] Fabian Gieseke, Justin Heinermann, Cosmin Oancea, and Christian Igel. Buffer kd trees: processing massive nearest neighbor queries on gpus. In *Proceedings of The 31st International Conference on Machine Learning*, pages 172–180, 2014.

[22] Merrill D. and Grimshaw A.S. High performance and scalable radix sorting: a case study of implementing dynamic parallelism for GPU computing. *Parallel Processing Letters*, 21(2):245–272, 2011.

A Additional details

In the following we review in greater detail the main complexities characterizing the various phases of the processing pipeline.

A.1 Index construction

In Section 2.2 we introduced the index construction phase. Here we detail the complexity characterizing this operation. The computation of single Morton codes has a fixed cost determined by the number of bits used for coordinate representation; therefore, computing Morton codes in parallel for all objects has a complexity equal to $O(|P|)$. Since Radix Sort [22] is used for the subsequent sorting operation, the complexity related to sorting is $O(b \cdot |P|)$, where b represents the base value used during sorting: given that $b \ll |P|$, the overall complexity can be approximated to $O(|P|)$.

After the sorting operation, the actual iterative PR-quadtree construction begins. The operation related to detecting quadrants which need to be split has a worst-case complexity equal to

$$O\left(l_{max} \cdot |P| + 2 \sum_{l=0}^{l_{max}} 4^l\right),$$

where the first term is due to the amount of objects scanned at each iteration, while $2 \cdot \sum_{l=0}^{l_{max}} 4^l = 2 \cdot \frac{1-4^{l_{max}-1}}{1-4}$ represents the maximum amount of starting and ending indices related to the 4^l quadtree quadrants at any level l. Since $4^l \ll |P|$, the related computational overhead is negligible and therefore the average complexity can be safely approximated to $O(l_{max} \cdot |P| + 2 \sum_{l=0}^{l_{max}} 4^l) \simeq O(l_{max} \cdot |P|)$. The remaining operation, carried on CPU, in which we determine which quadrants need to be split at the next level has a complexity equal to $2 \sum_{l=0}^{l_{max}} 4^l$ and, consequently, has negligible cost.

In conclusion, since we impose l_{max} to be a low constant (e.g., $l_{max} \leq 10$) and the amount of quadrants created per each level is orders of magnitude lower than the amount of objects, the overall complexity of the index creation phase can be approximated to $O(|P|)$.

A.2 Lookup table z_{map} initialization

In Section 2.2 we introduced the lookup table z_{map} used to retrieve in constant time the identifier of the quadtree leaf in which an object/query falls. Here we detail the complexity related to its initialization, which happens immediately after the index construction phase.

This operation is performed entirely on GPU by assigning each $\mathcal{C}$ cell (quadtree leaf) to a GPU streaming multiprocessor, which in turn initializes in parallel the interval of cells (elements of the lookup table) in $\mathcal{C}^{l_{deep}}$ contained by the $\mathcal{C}$ cell assigned. The overall complexity is therefore $O(|\mathcal{C}| + |\mathcal{C}^{l_{deep}}|)$.

A.3 Moving objects indexing

In Section 4.1.2 we introduced the *moving objects indexing* phase. Here we detail the related complexity.

The overall complexity of this subphase is $O(|P| + b \cdot |P| + |P| + 2|P| + |\hat{\mathcal{C}}|)$: the first term is related to the l_{deep} Morton codes parallel computation, while the second term is related to the subsequent sorting operation (by means of Radix Sort). The third term is due to the lookups in z_{map} needed to retrieve the final leaves identifiers. Finally, the fourth term is due to the double scan over the set of objects needed to detect the discontinuities between objects belonging to different cells, while $|\hat{\mathcal{C}}|$ represents the amount of active cells for which we actually have to write out the related indexing information. Since $|\hat{\mathcal{C}}| \ll |P|$, the overall complexity is in the order of $O(|P|)$.

A.4 Iterative k-NN queries computation

In the following we detail some of the complexities characterizing the main operations carried on during the *iterative k-NN queries computation* phase (Section 4.2).

A.4.1 First iteration

Query indexing and task materialization. This phase, which is described in Section 4.2.1, has a complexity which can be derived equivalently to the one characterizing the *moving objects indexing* phase, so the reader should refer to Appendix A.3). Consequently, the complexity of this phase is in the order of $O(|Q|)$.

Distance computations. In the following we describe the complexity of the *distance computations* phase presented in Section 4.2.1. Aside from specific patterns used to write out queries results, strategies based on caching or coalescing are, on the whole, almost identical. For this reason, we take Algorithm 1 as a point of reference. If we consider the operations carried on within each task, the complexity of the distance computation phase is mainly dictated by the scans over the set of objects belonging to each cell $c \in \overline{C}$, as well as by the k-selection algorithm. The overall number of scans inside Algorithm 1 are two, one at line 4 and the other one at lines 7-11; therefore, the related complexity is equal to $O(2 \cdot |\{P \cap c\}|)$. For what is related to the k-selection algorithm we have that the number of iterations strongly depends on local densities affecting the spatial distribution inside a cell. If d represents the minimum distance between pairs of objects in c and *numBins* represents the amount of bins used, such figure is equal to [10]:

$$O\left(\left\lceil\left|log_{numBins}(\frac{dist_{max} - dist_{min}}{d})\right|\right\rceil\right), \tag{1}$$

Summing up the considerations done above, if we denote the amount of iterations yielded by *findKDist* in the worst case as *maxIterations*, the overall complexity of the distance computation phase becomes:

$$O\Big((2 + maxIterations) \cdot |\{P \cap c\}|\Big). \tag{2}$$

A.4.2 Subsequent iterations

in the following we review the complexities of the main operation carried on during any subsequent iteration (Section 4.2.2).

Massive tree navigation. Query-wise, the complexity of this subphase is dictated by the set of active queries in the direction considered, i.e., $O(|Q_{processed}|)$; if we focus on a single query, the worst case scenario corresponds to visiting all quadtree nodes when such quadtree corresponds to a uniform grid, given that we use z_{map} to recover leaves identifiers. If l_{deep} is the deepest quadtree level, this corresponds to a complexity equal to $O(\sum_{l=0}^{l_{deep}} 4^l) = O(\frac{1-4^{l_{deep}-1}}{1-4})$. However, these complexities tell us little on the overall amount of iterations we have to expect from our approach, since this depends on portions of the tree involved by queries, which in turn depend on a complex mix of factors such as the objects spatial distribution, how such distribution may vary over the time, the quadtree height, how objects spread across the quadtree leaves and the amount of nearest neighbours per query to compute (k).

Query indexing and task materialization. The same reasoning done in Appendix A.4.1 applies here as well, however considering that the set of queries indexed at any iteration is limited to those active in the direction considered.

Distance computations. The overall complexity of this subphase can be determined by means of Equations 1 and 2, since the underlying algorithms are derived from the ones used in the *distance computation* phase performed during the first iteration (Section 4.2.1). More precisely, if c denotes the cell considered for a given query q, we have that the complexity related to the determination of the (up to) k nearest objects in c is, in the worst case

$$O\Big((2 + maxIterations) \cdot |\{p | p \in c \wedge d(p,q) < MAXDIST_q\}|\Big),$$

while the operations needed to fuse the set of (up to) k nearest objects in c with the (up to) k nearest objects in the query result set computed so far is, in the worst case,

$$O\Big((2 + maxIterations) \cdot 2k\Big).$$